Your response

Question	Your response
Question 1: Ofcom's general approach to information gathering (Section 3 of the draft guidance) **Do you have any comments on Ofcom's proposed general approach to information gathering, as outlined in Section 3 of the draft guidance?**	Confidential? – N The new information gathering powers bestowed upon Ofcom by the Online Safety Act mark a significant step forward in holding service providers to account. The rationale for these statutory powers outlined in §3.1-3.6 is welcome: we strongly concur with the need to ensure that "regulatory decisions are founded on a robust evidence base", and that without these powers this can be impeded by the "information asymmetry that may exist between Ofcom and stakeholders". Overcoming these external access challenges is key - whilst these new powers and the accompanying guidance are necessary for doing this, we emphasise that in practice the mechanism by which external access to proprietary systems is provided will have a significant impact on how effectively Ofcom can exercise these new powers. §3.13 highlights factors Ofcom will consider prior to exercising its information gathering powers, including the associated resource costs, privacy concerns, and intellectual property risks. We agree that these are legitimate considerations, but highlight that improper understanding of their impact could inadvertently restrict the speed and rigour with which tests and audits can be carried out, which may mean risks and harms go undetected. For the past six years, OpenMined has worked on developing freely available open-source infrastructure that seeks to address these challenges, by facilitating external access to proprietary data and AI systems. This infrastructure can enable access whilst protecting privacy, security, and IP by design through the use of novel privacy enhancing technologies. This infrastructure can thus help operationalise the privacy and confidentiality requirements set out in §3.13 and §3.19. Furthermore, this infrastructure can enable Ofcom (and appointed skilled-persons) to securely audit or execute tests *remotely,* minimising the need to go on-site, and thus minimising the associated costs. Our response to subsequent questions will aim to describe how this infrastructure works in practice, and

Question	Your response
	situate it in the context of the specific information gathering powers described in the guidance.
Question 2: Information notices (Section 4 of the draft guidance) **a) Information notices** **Do you have any comments on Ofcom's proposed approach to the process for issuing and responding to information notices.** **b) Requiring a test** **Do you have any comments on our proposed approach to information notices that require recipients to perform a test?** **c) Remote viewing** **Do you have any comments on our proposed approach to Remote Viewing Information Notices? For example, to the factors that we may take into account when considering whether to issue a Remote Viewing Information Notice.** **d) Coroner Information Notices** **Do you have any comments on our proposed approach to issuing Coroner Information Notices for the purpose of responding to requests for information by investigating authorities in connection with an investigation or inquest into the death of a child?** **e) Naming a senior manager** **Do you have any comments on the section relating to naming a senior manager who is in a position to ensure compliance with an information notice?**	Confidential? – N Our response to this question will focus on mechanisms for operationalising Ofcom's powers of remote viewing and requiring a test. This response will also incorporate aspects of Question 3, as mechanisms for requiring tests or remote viewing by Ofcom can be readily extended to appointed skilled-persons. The mechanism by which external access to proprietary systems is provided can have a significant impact on how effectively Ofcom (and appointed skilled persons) can exercise their new powers. Traditional approaches to external access require that a representative from an oversight organisation either: 1) Travels on-site to have direct access to the system, or 2) Relies on the service provider to build bespoke APIs that facilitate the proposed assessment or evaluation. These legacy approaches can introduce significant delays, with 1) placing significant restrictions on where and when audits can be carried out, and the level of vetting required; and 2) requiring several months of development time to create a new API. As well as introducing delays, these approaches also introduce additional costs for both the service provider and the oversight organisation, such as the travel costs for going on-site, or the development costs of creating new APIs (e.g., expensive engineers and product manager's time). Based on our experience facilitating external access with service providers including LinkedIn and Reddit, we know better solutions are possible. The legacy approaches above — which can take months to facilitate very narrow access — can be transferred into an hours-to-days long process using a new class of oversight tools. 1. Remote Execution In this oversight setup, a service provider can load the assets (e.g., user logs, impression data, etc.) necessary to

Question	Your response
	evaluate compliance with a specific duty into a high-side server deployed on their local infrastructure inside their firewall. They would then deploy a low-side server that contains mock assets—assets that directly imitate the structure of the real assets but contain fake, non-sensitive information. The low-side server would be shared with oversight organisations so that they could prepare, test, and iterate on their audit/evaluation code using mock assets downloaded to their local machine. This step ensures that oversight organisations specify their code with appropriate precision to get the appropriate result. We refer to the high-side and low-side servers together as comprising the service provider's datasite. Once content with their code, the oversight organisation can share it with the service provider. The service provider confirms the audit/evaluation goals are as specified in the code and can then approve the project to be executed against the private assets on the high-side server and return the result to the oversight organisation. The oversight organisation now has the results of an audit/evaluation run against the private assets, crucially without ever directly seeing the assets. 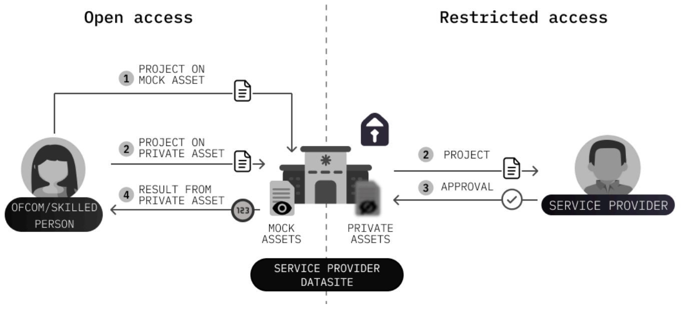 Open access Restricted access 1 PROJECT ON MOCK ASSET 2 PROJECT ON PRIVATE ASSET 4 RESULT FROM PRIVATE ASSET 123 OFCOM/SKILLED PERSON MOCK ASSETS PRIVATE ASSETS SERVICE PROVIDER DATASITE 2 PROJECT 3 APPROVAL SERVICE PROVIDER This tool mitigates the legitimate barriers to access that have thwarted legacy oversight tools. However, it still relies on a degree of trust between the service provider and the oversight organisation in that the oversight organisation trusts the service provider to move their code from the low-side server to the high-side server and run it as it is. Although a degree of trust is still required in this new setup, it is significantly less than in current oversight setups. This approach could streamline the process for enabling an independent third-party to carry out a skilled-persons'

Question	Your response
	assessment, as well as enabling Ofcom to develop and see the results of an empirical test executed against test data provided by the service provider. This can be done *without requiring the third-party to travel on-site.* As well as minimising costs, this also provides a practical way for Ofcom to be able to remotely view or carry out tests against a service provider “even in circumstances where the recipient does not have premises in the United Kingdom”, as per §4.52. 2. Secure Enclaves Building on the remote execution tools described above, an oversight organisation may want to run an audit/evaluation that contains classified or sensitive information that is not appropriate or legal to share with a service provider. The oversight organisation can still use the low-side server to develop and refine its code, but now requires a different execution setup. In this situation, an oversight setup leveraging secure enclaves for mutual secrecy can be beneficial. Upon agreement to conduct the audit/evaluation, a secure enclave can be spun up on either the service provider's infrastructure, the oversight organisation’s infrastructure, or a trusted third-party’s infrastructure. The service provider can then send the required assets into the secure enclave, and the oversight organisation can send their code to be run. In this setup, the oversight organisation could also send their own test dataset to the enclave, enabling evaluations to be run against this dataset whilst keeping it secret from the service provider. When the audit/evaluation is complete, the results can be returned to the specified parties (i.e., just the oversight organisation or both the oversight organisation and the service provider). 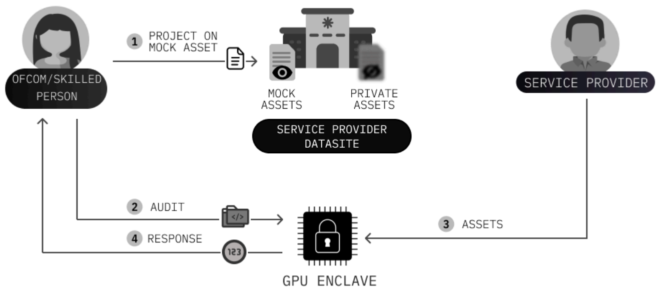 1 PROJECT ON MOCK ASSET OFCOM/SKILLED PERSON MOCK ASSETS PRIVATE ASSETS SERVICE PROVIDER DATASITE SERVICE PROVIDER 2 AUDIT 4 RESPONSE 3 ASSETS GPU ENCLAVE

Question	Your response
	This tool mitigates the legitimate barriers to access that have thwarted legacy oversight tools and again reduces the level to which an oversight organisation must trust the service provider. However, it still relies on a degree of trust between the service provider and the oversight organisation in that the oversight organisation must trust the service provider to send the expected assets necessary for the audit/evaluation into the secure enclave. Ofcom could leverage this setup in situations where it may be necessary to run an empirical test against a proprietary system using a test dataset provided by Ofcom (as per §4.40), but which Ofcom needs to keep secret from the service provider. Even if the test dataset does not contain sensitive information, it may still be beneficial to keep it secret from the service provider in order to prevent the common problem of the service provider configuring their system to overfit to the test dataset, which can produce metrics that are not reflective of the system’s true performance. The end-to-end system needed to realise the full vision of this oversight setup does not exist yet, but most of the individual pieces of technology necessary to power such a system do exist and are ready for pre-deployment testing. Ofcom could pilot such a system in collaboration with relevant science and technology departments (e.g., DSIT, the ONS) to finalise and mature this technology. 3. Public Proof Registry This oversight setup extends parts of the previous setup by leveraging the cryptographic hashes provided by secure enclaves to create a robust chain of custody from the service provider to the regulator and, ultimately, to the end user. In this setup, when an oversight organisation wishes to perform an audit/evaluation, the service provider loads the relevant assets into a secure enclave and adds a public proof to a registry (through Zero-Knowledge Proofs or other cryptographic methods). This registry could be hosted by Ofcom. The registry could then be made available to any independent third-party appointed by Ofcom as a skilled-person to use as a check when they audit, evaluate, assess, or

<table>
<tr><th>Question</th><th>Your response</th></tr>
<tr><td></td><td>
otherwise inspect a system that they are verifiably working with the assets they expect.

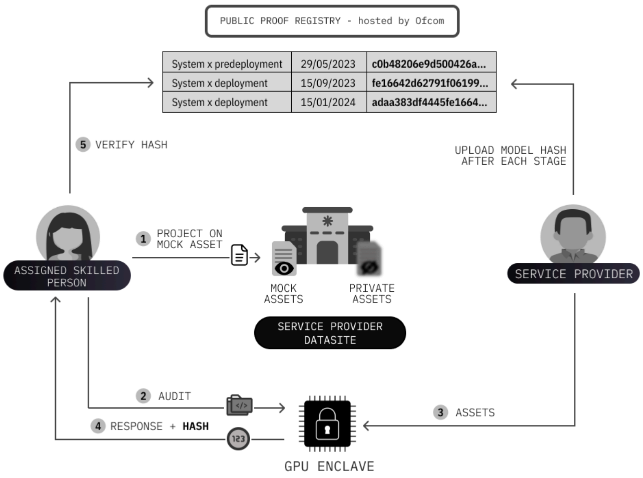

<table>
<tr><td>System x predeployment</td><td>29/05/2023</td><td>c0b48206e9d500426a...</td></tr>
<tr><td>System x deployment</td><td>15/09/2023</td><td>fe16642d62791f06199...</td></tr>
<tr><td>System x deployment</td><td>15/01/2024</td><td>adaa383df4445fe1664...</td></tr>
</table>

As with the oversight setups above, this setup also mitigates the legitimate barriers to access that have thwarted legacy oversight tools, but the additional tool in this setup ultimately eliminates the need for oversight organisations to trust the service provider since they can independently verify the claims.

§4.42 permits service providers to carry out empirical tests in separate ‘test environments’. Whilst there are circumstances where it may be sensible to separate testing from the live production environment, this introduces a risk that the system that is under test differs from the system that is in production (and they could drift further apart over time). The use of a public proof registry could mitigate this risk, by providing a secure mechanism by which Ofcom (or another authorised third-party) can verify whether software running in production is the same version that was tested against.

Aside from the enclave dependencies, kickstarting a public proof registry can be done relatively quickly; it requires participation from service providers to submit hashes and commitment from a trusted party (which could be Ofcom) to host and maintain the registry.
</td></tr>
<tr><td><b>Question 3: Skilled persons’ reports (Section 5 of the draft guidance)</b></td><td>Confidential? – N</td></tr>
</table>

Question	Your response
Do you have any comments on our approach to skilled persons' reports? This might include when we might decide to require a skilled person's report, and the typical process that we propose to follow.	Please see our response to Question 2, which describes a streamlined process and mechanism for enabling a skilled person to effectively carry out an assessment.
Question 4: Interviews (Section 6 of the draft guidance) **Do you have any comments on the section of guidance dealing with the power to require an individual to attend an interview?**	Confidential? – Y / N
Question 5: Entry with or without a warrant (Section 7 of the draft guidance) **Do you have any comments on our proposed approach to entry either with or without a warrant? This might include the typical process and our interpretation of the requirement to have regard to the Home Office's code of practice on powers of entry.**	Confidential? – Y / N
Question 6: Audit (Section 7 of the draft guidance) **Do you have any comments on our proposed approach to the power for Ofcom to carry out an audit to assess compliance?**	Confidential? – N Whilst we understand that powers of entry and inspection should typically be reserved for more complex cases, we challenge the implication in §7.2 that the same threshold should apply for audit of services. §7.44 provides a broad set of actions Ofcom can take under an audit notice, some of which may require this higher bar to be met (e.g. where the audit itself requires entry to a premises). However, for software systems, effective external audit can (and arguably should) be carried out much more frequently, utilising the remote execution infrastructure described in our response to Question 2. Without this, such a high bar is likely to mean that in practice audits are only carried out after a system has already led to significant harm. A more proactive

Question	Your response
	auditing regime can help ensure that these harms do not occur in the first place.
Question 7: Consequences of failure to comply with an information power (Section 8 of the draft guidance) **Do you have any comments on the potential consequences of a failure to comply with any of the information gathering powers covered in the draft guidance? This might be either on breaches that may be subject to enforcement action by Ofcom, or those that may constitute criminal offences.**	Confidential? – Y / N
Question 8: Additional comments **Do you have any other comments on the draft guidance?** **Please provide any information or evidence in support of your views.**	Confidential? – Y / N

Please complete this form in full and return to OSinfoguidance@ofcom.org.uk